Your search matched **0** of **1071730** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

(flush buffer <sentence> (store in cache)) and (read onl Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

Results:
No documents matched your query.

Welcome
United States Patent and Trademark Office

» Advanced Searc

Help FAQ Terms IEEE Peer Review Quick Links

Welcome to IEEE Xplore®
- Home
- What Can I Access?
- Log-out

Tables of Contents
- Journals & Magazines
- Conference Proceedings
- Standards

Search
- By Author
- Basic
- Advanced

Member Services
- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise
- Access the IEEE Enterprise File Cabinet

Try our New Full-text Search Prototype GO Help

1) Enter a single keyword, phrase, or Boolean expression. Example: acoustic imaging (means the phrase acoustic imaging plus any stem variations)
2) Limit your search by using search operators and field codes, if desired.
Example: optical <and> (fiber <or> fibre) <in> ti
3) Limit the results by selecting Search Options.
4) Click Search. See Search Examples

```
(flush buffer <sentence>
(store-in cache)) and (read-
only <near/4> cache) and
(store-through <near/4>
```

Start Search Clear

Note: This function returns plural and suffixed forms of the keyword(s).

Search operators: <and> <or> <not> <in> More

Field codes: au (author), ti (title), ab (abstract), jn (publication name), de (index term) More

Search Options:

Select publication types:

- [x] IEEE Journals
- [x] IEE Journals
- [x] IEEE Conference proceedings
- [x] IEE Conference proceedings
- [x] IEEE Standards

Select years to search:

From year: All to Present

Organize search results by:

Sort by: Relevance
In: Descending order

List 15 Results per page

Terms used
flush buffer sentence store in **cache** and **read only near/4 cache** and **store through near/4 cache**

Found **10,368** of **142,346**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Efficient (stack) algorithms for analysis of write-back and sector memories

James G. Thompson, Alan Jay Smith

January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Full text available: pdf(2.93 MB) Additional Information: full citation, abstract, references, citings, index terms, review

For the class of replacement algorithms known as stack algorithms, existing analysis techniques permit the computation of memory miss ratios for all memory sizes simultaneously in one pass over a memory reference string. We extend the class of computations possible by this methodology in two ways. First, we show how to compute the effects of copy-backs in write-back caches. The key observation here is that a given block is clean for all memory sizes less than or equal to C ...

2 Issues related to MIMD shared-memory computers: the NYU ultracomputer approach

Jan Edler, Allan Gottlieb, Clyde P. Kruskal, Kevin P. McAuliffe, Larry Rudolph, Marc Snir, Patricia J. Teller, James Wilson

June 1985 **ACM SIGARCH Computer Architecture News , Proceedings of the 12th annual international symposium on Computer architecture**, Volume 13 Issue 3

Full text available: pdf(998.36 KB) Additional Information: full citation, citings, index terms

3 Cache Memories

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: pdf(4.61 MB) Additional Information: full citation, references, citings, index terms

4 Garbage collection for a client-server persistent object store

Laurent Amsaleg, Michael J. Franklin, Olivier Gruber

August 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 3

Full text available: pdf(267.18 KB) Additional Information: full citation, abstract, references, citings, index terms, review

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environmnet. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does non require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

5 Code optimization - I: Local scheduling techniques for memory coherence in a clustered VLIW processor with a distributed data cache

Enric Gibert, Jesús Sánchez, Antonio González

March 2003 **Proceedings of the international symposium on Code generation and optimization: feedback-directed and runtime optimization**

Full text available: pdf(1.19 MB) Additional Information: full citation, abstract, references, index terms

Publisher Site

Clustering is a common technique to deal with wire delays. Fully-distributed architectures, where the register file, the functional units and the cache memory are partitioned, are particularly effective to deal with these constraints and besides they are very scalable. However, the distribution of the data cache introduces a new problem: memory instructions may reach the cache in an order different to the sequential program order, thus possibly violating its contents. In this paper two local sch ...

6 Cache write policies and performance

Norman P. Jouppi

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture**, Volume 21 Issue 2

Full text available: pdf(1.14 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper investigates issues involving writes and caches. First, tradeoffs on writes that miss in the cache are investigated. In particular, whether the missed cache block is fetched on a write miss, whether the missed cache block is allocated in the cache, and whether the cache line is written before hit or miss is known are considered. Depending on the combination of these polices chosen, the entire cache miss rate can vary by a factor of two on some applications. The combination of no- ...

7 Consistency management for virtually indexed caches

Bob Wheeler, Brian N. Bershad

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems**, Volume 27 Issue 9

Full text available: pdf(1.62 MB) Additional Information: full citation, references, citings, index terms

8 ENWRICH: a compute-processor write caching scheme for parallel file systems

Apratim Purakayastha, Carla Schlatter Ellis, David Kotz

May 1996 **Proceedings of the fourth workshop on I/O in parallel and distributed systems: part of the federated computing research conference**

Full text available: pdf(1.38 MB) Additional Information: full citation, references, citings, index terms

9 The Rio file cache: surviving operating system crashes

Peter M. Chen, Wee Teck Ng, Subhachandra Chandra, Christopher Aycock, Gurushankar Rajamani, David Lowell

September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings, index terms

One of the fundamental limits to high-performance, high-reliability file systems is memory's vulnerability to system crashes. Because memory is viewed as unsafe, systems periodically write data back to disk. The extra disk traffic lowers performance, and the delay period before data is safe lowers reliability. The goal of the Rio (RAM I/O) file cache is to make

ordinary main memory safe for persistent storage by enabling memory to survive operating system crashes. Reliable memory enables a syste ...

10 Cache memory performance in a unix enviroment

Cedell Alexander, William Keshlear, Furrokh Cooper, Faye Briggs

June 1986 **ACM SIGARCH Computer Architecture News**, Volume 14 Issue 3

Full text available: pdf(2.10 MB) Additional Information: full citation, citings, index terms

11 Silent stores for free

Kevin M. Lepak, Mikko H. Lipasti

December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available: pdf(521.75 KB) ps(1.97 MB) Additional Information: full citation, references, citings, index terms

Publisher Site

12 The directory-based cache coherence protocol for the DASH multiprocessor

Daniel Lenoski, James Laudon, Kourosh Gharachorloo, Anoop Gupta, John Hennessy

May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, citings, index terms

DASH is a scalable shared-memory multiprocessor currently being developed at Stanford's Computer Systems Laboratory. The architecture consists of powerful processing nodes, each with a portion of the shared-memory, connected to a scalable interconnection network. A key feature of DASH is its distributed directory-based cache coherence protocol. Unlike traditional snoopy coherence protocols, the DASH protocol does not rely on broadcast; instead it uses point-to-point messages sent between th ...

13 An economical solution to the cache coherence problem

James Archibald, Jean Loup Baer

January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3

Full text available: pdf(728.73 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we review and qualitatively evaluate schemes to maintain cache coherence in tightly-coupled multiprocessor systems. This leads us to propose a more economical (hardware-wise), expandable and modular variation of the "global directory" approach. Protocols for this solution are described. Performance evaluation studies indicate the limits (number of processors, level of sharing) within which this approach is viable.

14 Supporting reference and dirty bits in SPUR's virtual address cache

D. A. Wood, R. H. Katz

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture**, Volume 17 Issue 3

Full text available: pdf(1.12 MB) Additional Information: full citation, abstract, references, citings, index terms

Virtual address caches can provide faster access times than physical address caches, because translation is only required on cache misses. However, because we don't check the translation information on each cache access, maintaining reference and dirty bits is more difficult. In this paper we examine the trade-offs in supporting reference and dirty bits in a virtual address cache. We use measurements from a uniprocessor SPUR prototype to evaluate different alternatives. The prototype's buil ...

15

Informed prefetching and caching
R. H. Patterson, G. A. Gibson, E. Ginting, D. Stodolsky, J. Zelenka
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5
Full text available: pdf(2.13 MB) Additional Information: full citation, references, citings, index terms

16 Compiler and hardware support for cache coherence in large-scale multiprocessors: design considerations and performance study
Lynn Choi, Pen-Chung Yew
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2
Full text available: pdf(1.48 MB) Additional Information: full citation, abstract, references, index terms

In this paper, we study a hardware-supported, compiler directed (HSCD) cache coherence scheme, which can be implemented on a large-scale multiprocessor using off-the-shelf microprocessors, such as the Cray T3D. It can be adapted to various cache organizations, including multi-word cache lines and byte-addressable architectures. Several system related issues, including critical sections, inter-thread communication, and task migration have also been addressed. The cost of the required hardware sup ...

17 Software-controlled caches in the VMP multiprocessor
D. R. Cheriton, G. A. Slavenburg, P. D. Boyle
June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture**, Volume 14 Issue 2
Full text available: pdf(958.63 KB) Additional Information: full citation, abstract, references, citings, index terms

VMP is an experimental multiprocessor that follows the familiar basic design of multiple processors, each with a cache, connected by a shared bus to global memory. Each processor has a synchronous, virtually addressed, single master connection to its cache, providing very high memory bandwidth. An unusually large cache page size and fast sequential memory copy hardware make it feasible for cache misses to be handled in software, analogously to the handling of virtual memory page faults. Har ...

18 Optimizing the data cache performance of a software MPEG-2 video decoder
Peter Soderquist, Miriam Leeser
November 1997 **Proceedings of the fifth ACM international conference on Multimedia**
Full text available: pdf(1.76 MB) Additional Information: full citation, references, citings, index terms

19 Implementing a cache for a high-performance GaAs microprocessor
O. A. Olukotun, T. N. Mudge, R. B. Brown
April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture**, Volume 19 Issue 3
Full text available: pdf(1.12 MB) Additional Information: full citation, references, citings, index terms

20 Characterizing the Storage Process and Its Effect on the Update of Main Memory by Write Through
Alan Jay Smith
January 1979 **Journal of the ACM (JACM)**, Volume 26 Issue 1
Full text available: pdf(1.27 MB) Additional Information: full citation, references, citings, index terms